Scheduling using Interactive Oracles: Connection between Iterative Optimization and Low-complexity Scheduling

Jinwoo Shin
Department of Electrical Engineering
Korea Advanced Institute of Science and Technology (KAIST)
Republic of Korea
jinwoos@kaist.ac.kr

Tonghoon Suk
H. Milton Stewart School of Industrial & Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA
tonghoons@gatech.edu

arXiv:1407.3694v1 [math.PR] 14 Jul 2014

ABSTRACT

Since Tassiulas and Ephremides proposed the maximum weight scheduling algorithm of throughput-optimality for constrained queueing networks in 1992, extensive research efforts have been made for resolving its high complexity issue under various directions. In this paper, we resolve the issue by developing a generic framework for designing throughput-optimal and low-complexity scheduling algorithms. Under the framework, an algorithm updates current schedules via an interaction with a given oracle system that can generate a solution of a certain discrete optimization problem in a finite number of interactive queries. Therefore, one can design a variety of scheduling algorithms under this framework by choosing different oracles, e.g., the exhaustive search (ES), the markov chain monte carlo (MCMC), the belief propagation (BP) and the cutting-plane (CP) algorithms. The complexity of the resulting algorithm is decided by the number of operations required for an oracle processing a single query, which is typically very small. Somewhat surprisingly, we prove that an algorithm using any such oracle is throughput-optimal for general constrained queueing network models that arise in the context of emerging large-scale communication networks. In particular, the 'pick-and-compare' algorithms developed by Tassiulas in 1998 and recently developed queue-based CSMA algorithms can be also understood as special cases of such algorithms using ES and MCMC oracles, respectively. To our best knowledge, our result is the first that establishes a rigorous connection between iterative optimization methods and low-complexity scheduling algorithms, which we believe provides various future directions and new insights in both areas.

1. INTRODUCTION

The problem of dynamic resource allocation in constrained queueing networks is fundamental to modern communication networks. This has long been the subject of significant research, where the seminal work [1] of Tassiulas and Ephremides proposed the MW (maximum weight) scheduling algorithm of throughput-optimality. It has been also a powerful paradigm in cross-layer control and resource allocation [2]. However, the algorithm cannot be implemented in practice due to its high algorithmic complexity. Much effort has been made since then for developing algorithms of lower complexity while maintaining the same performance guarantee. Examples of such algorithms include greedy algorithms [3, 4, 5, 6] and random access algorithms [7, 8, 9].

Recently, there have been exciting progresses for the topic that even fully distributed random access algorithms based on CSMA (Carrier Sense Multiple Access) can achieve the desired high performance, i.e., throughput-optimality, in wireless interference network models, e.g., see [10, 11, 12]. The main intuition underlying these results is that nodes in a network can adjust their CSMA parameters dynamically using local information such as queue-length so that they can simulate the MW algorithm asymptotically for throughput-optimality. However, it is known that the queue-based CSMA algorithms suffer from very poor delay performance. There have been also significant efforts formally addressing the limitation in the delay performance of the queue-based CSMA algorithms [13, 14] as well as designing better alternatives in delay under certain scenarios [15, 16, 17, 18]. In addition, another issue is that the throughput-optimality proof in [10, 12] is very sensitive to the underlying wireless network model. Namely, it is not an easy task to design similar CSMA-like throughput-optimal algorithms in other network models since one has to reproduce painful, long and routine proof-arguments for throughput-optimality in each different model. The main goal of this paper is to resolve these issues of the recent queue-based CSMA algorithms by developing a general framework that can produce a variety of throughput-optimal and low-complexity scheduling algorithms, not sensitive to the underlying network models.

1.1 Related Work: MW and Variants

Design and analysis of resource allocation algorithms in constrained queueing networks have been of great interest for last four decades across various research communities, where the contention resolution protocols in multiple-access broadcast channels have been studied since the mid-1970s [19]. The broadcast channels induce a special constraint in resources usages among queues (i.e., every queue contend with all other queues), and simple distributed contention resolution protocols gained much attention in the ages. It

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...$15.00.

is the year 1992 that the seminal work by Tassiulas and Ephremides proposed a throughput-optimal algorithm, referred to as MW, for general resource constraints. Its main appeal is that it does require only information on the current queue lengths (and instantaneous service rates) and does not rely on the knowledge of underlying system parameters. However, as we mentioned earlier, the MW algorithm suffers from the high implementation complexity, and a huge array of subsequent research has been made to develop algorithms with high performance guarantees and low-complexity.

Maximal scheduling or Longest-Queue-First algorithm are low-complexity alternatives to MW, but they achieve only some fraction of the maximal throughput region. Parallel Iterative Matching [3] and iSLIP [4] were shown to be 50% throughput-optimal [20]. On the other hand, Kumar et al. [21] and Dimakis and Walrand [25] identified sufficient conditions on the network topology for throughput-optimality. Those conditions were further analyzed to obtain fractional throughput results about a class of wireless networks by Joo et al. [5] and Leconte et al. [6]. These algorithms are generally not throughput-optimal and require multiple rounds of message exchanges between nodes.

Simpler or distributed implementations of MW has been also extensively studied in the literature. Tassiulas [22] provides linear-complexity versions of MW, called 'pick-and-compare', but it is known that they suffer from the bad delay performance. The works on this line also include a variant of MW by Giaccone, Prabhakar and Shah [23] and a gossip-based algorithm by Modiano, Shah and Zussman [24]. However, these algorithms are specific to certain network models and still require high information (or message) exchanges for each new scheduling decision. A few years ago, it was proven that even fully distributed queue-based random access algorithms based on CSMA with very little message passing can simulate the MW algorithm asymptotically and hence achieve throughput-optimality in wireless interference network models [10, 12].

However, the CSMA algorithms are known to suffer from the following two issues: the bad delay performance and the limited applicability to other network models. For the delay issue, Bouman et al. [14] formally established explicit lower bounds on the delay of queue-based random access algorithms. On the other hand, significant efforts have also been made to design CSMA algorithms with better delay performance, e.g., see [15, 16, 17, 18]. The limited applicability issue is rather technical: it is a technically tedious task to design similar CSMA-like throughput-optimal algorithms in other network models since one has to reproduce painful, long and routine proof-arguments for analyzing 'time-varying' Markov chains in each different model, e.g., see the buffered circuit switched network model [12]. The main motivation of this paper is for resolving these issues and goes beyond CSMA-like algorithms for achieving high performance guarantees and low implementation complexity, not sensitive to the underlying network models.

1.2 Contribution: Beyond CSMA

We develop a generic framework for designing throughput-optimal and low-complexity scheduling algorithms, where queue-based CSMA algorithms belong to it as special cases. Under the framework, an algorithm updates current schedules via an interaction with a given oracle system that can generate a solution of a certain discrete optimization problem in a finite number of interactive queries. Specifically, an algorithm maintains three parameters, 'schedule', 'weight' and 'advice', at each time instance. The schedule and advice are kept being updated under interactions with the oracle system and the weight is decided by some function of current queue-size, which we call 'queue-length-function'. Therefore, the scheduling algorithm design is determined by 'which oracle is used' and 'which queue-length-function is chosen'. The oracle system can be designed through iterative optimization methods including the naive exhaustive search algorithm, the markov chain monte carlo (e.g., the glauber dynamics [26]), the message-passing inference algorithms (e.g., the belief propagation [27]) and the linear/integer programming (e.g., the cutting-plane method [28]). In this paper, we describe three oracle examples, the exhaustive search (ES), the markov chain monte carlo (MCMC) and the belief propagation (BP) in Section 4. In particular, the 'pick-and-compare' algorithms [22] and queue-based CSMA algorithms can be understood as such algorithms (with minor modifications) using ES and MCMC oracles, respectively (see Section 4). The scheduling complexity is decided by the number of operations required for an oracle processing a single query, which is typically very small, e.g., ES, MCMC and BP.

We prove that given any oracle system, there exists a queue-length-function such that an algorithm using the oracle is throughput-optimal under general constrained queueing network models, which includes wireless network models, input-queued switch models, the buffered circuit switched network model, etc. In particular, we first show there exists a queue-length-function such that its value does not change too often (i.e., remains fixed for a long time) and hence the MW optimization problem can be solved through (enough) interactions to the oracle, i.e., the algorithm simulates MW and is throughput-optimal. This proof strategy is similar to those of [10, 12] in spirit. The main difference is that we force the queue-length-function value being 'rarely-varying' in integers, while the authors in [10, 12] do it being 'slowly-varying' in real numbers. Due to such a difference, now we do not need to analyze 'time-varying' algorithms (e.g., Markov chains), i.e., analysis when algorithm parameters are fixed is enough. This makes the proof much simpler and not sensitive to the given oracle and underlying network model.

Our strong theoretical statement provides a quite powerful paradigm for developing a variety of low-complexity scheduling algorithms of high performance. For example, our experimental results show that the delay performance of such an algorithm using the BP oracle is close to that of the MW algorithm, and it significantly outperforms CSMA algorithms (see Section 6). Moreover, one can also easily design other types of throughput-optimal algorithms, via designing various oracle systems, under a quite broad class of constrained queueing network models, without suffering from repeating tedious proofs sensitive to the underlying assumptions and setups. At a high level, we use the interactive oracle concept to provide a 'black-box' tool for designing scheduling algorithms of high performance and low-complexity. This establishes a rigorous connection between iterative optimization methods and low-complexity scheduling algorithms, which we believe provides various future directions and new insights in both areas.

1.3 Organization

In Section 2, we describe the constrained queueing network model of our interest and the performance metric (i.e., throughput-optimality) for scheduling algorithms. Section 3 provides the main result of this paper, i.e., our generic framework for designing a throughput-optimal and low-complexity scheduling algorithm which decides its current schedule via an interaction with an oracle system. In Section 4, we provide several examples of scheduling algorithms under the framework. Section 5 provides the formal proof of the main result, and our experimental results are presented in Section 6.

2. PRELIMINARIES

2.1 Network Model

We consider a stochastic network of n queues represented by V with $|V| = n$, and time is discrete indexed by $t \in \mathbb{Z}_+ := \{0, 1, \dots\}$. External packets (or jobs) arrive at each queue as per an exogenous stochastic process, where we let $A_i(t)$ denote the number of packets arriving at queue i during the time interval $[t, t+1)$. We assume that $\{A_i(t) \in \{0,1\} : t \in \mathbb{Z}_+, i \in V\}$ are independent Bernoulli random variables with

$$\mathbb{E}[A_i(t)] = \lambda_i$$

for some $\lambda_i \in [0,1]$, called the *arrival rate* for queue i.

Packets in queues are served at unit rate subject to scheduling constraints. Formally, a schedule is denoted by $\boldsymbol{\sigma} = [\sigma_i : i \in V] \in \{0,1\}^n$, where $\sigma_i = 1$ means queue i is transmitting (or active) and $\sigma_i = 0$ otherwise. We let $S \subset \{0,1\}^n$ be a set of all available schedules and let $\boldsymbol{\sigma}(t) \in S$ be the schedule during $[t, t+1)$ for $t \in \mathbb{Z}_+$.

To describe the dynamics in the network, we let $Q_i(t) \in \mathbb{Z}_+$ denote the number of packets in i-th queue at time t (i.e., *length of queue*) and let $\boldsymbol{Q}(t) = [Q_i(t)]$ be the vector of queue lengths. Then the queueing dynamics for queue i can be described as

$$Q_i(t+1) \;=\; Q_i(t) + A_i(t) - \sigma_i(t) I_{\{Q_i(t)>0\}},$$

where $t \in \mathbb{Z}_+$ and $I_{\{x\}} = \begin{cases} 1 & \text{if } x \text{ is 'true'} \\ 0 & \text{otherwise} \end{cases}$.

For given V and S, our goal is to develop low-complexity and high-performance scheduling algorithms which determine $\boldsymbol{\sigma}(t)$ for each time instance $t \in \mathbb{Z}_+$.

2.2 Performance Metric

A scheduling algorithm is required to choose $\boldsymbol{\sigma}(t) \in S$ at each time instance $t \in \mathbb{Z}_+$. We say an algorithm is throughput-optimal if it keeps queues finite for as large a set of arrival rates $\boldsymbol{\lambda} = [\lambda_i : i \in V]$ as possible. To describe it formally, we define the capacity region as follows:

$$\mathcal{C} := \left\{ \sum_{\boldsymbol{\sigma} \in S} \alpha_{\boldsymbol{\sigma}} \boldsymbol{\sigma} \;:\; \sum_{\boldsymbol{\sigma} \in S} \alpha_{\boldsymbol{\sigma}} = 1 \text{ and } \alpha_{\boldsymbol{\sigma}} \geq 0 \text{ for all } \boldsymbol{\sigma} \in S \right\},$$

where it is the convex hull of the feasible scheduling set S. Since $\boldsymbol{\sigma}(t) \in S$, the effective 'service' rate induced by any algorithm over time is essentially in $\mathcal{C}$. Therefore, if arrival rates $\boldsymbol{\lambda}$ can be served by any algorithm, then there exists $\boldsymbol{\sigma} \in \mathcal{C}$ such that $\boldsymbol{\lambda} \leq \boldsymbol{\sigma}$ component-wise. We say $\boldsymbol{\lambda}$ is *(strictly) admissible* if λ is (strictly) less than $\boldsymbol{\sigma}$ for some $\boldsymbol{\sigma} \in \mathcal{C}$. Let Λ and Λ^o be the set of all admissible and strictly admissible rates, respectively:

$$\begin{aligned} \Lambda &:= \{\boldsymbol{\lambda} \in \mathbb{R}_+^n \;:\; \boldsymbol{\lambda} \leq \boldsymbol{\sigma}, \text{ for some } \boldsymbol{\sigma} \in \mathcal{C}\} \\ \Lambda^o &:= \{\boldsymbol{\lambda} \in \mathbb{R}_+^n \;:\; \boldsymbol{\lambda} < \boldsymbol{\sigma}, \text{ for some } \boldsymbol{\sigma} \in \mathcal{C}\}. \end{aligned}$$

A throughput-optimal algorithm keeps queues finite for any $\boldsymbol{\lambda} \in \Lambda^o$, where there can be many different ways to argue the finiteness of queues. In this paper, we use the following popular definition of throughput-optimality.

DEFINITION 1. *A scheduling algorithm is called* throughput-optimal *if for any $\boldsymbol{\lambda} \in \Lambda^o$ the (appropriately defined) network Markov chain is positive recurrent.*

If the state of the network Markov chain includes queue-sizes, then its positive recurrence implies that the queues remain finite with probability 1, which motivates the above definition.

2.3 Lyapunov-Foster Criteria

We recall the definition of the positive recurrence of a Markov chain $\{\boldsymbol{X}(t) : t \in \mathbb{Z}_+\}$ on a countable state space Ω. A state $\mathsf{x} \in \Omega$ is said to be *recurrent* if $\mathbb{P}[\tau_{\mathsf{x}} < \infty \,|\, \boldsymbol{X}(0) = \mathsf{x}] = 1$, where $\tau_{\mathsf{x}} = \inf\{t \geq 1 : \boldsymbol{X}(t) = \mathsf{x}\}$. If $\mathbb{E}[\tau_{\mathsf{x}} \,|\, \boldsymbol{X}(0) = \mathsf{x}] < \infty$, a recurrent state x is called *positive recurrent*. If a Markov chain is irreducible and it has a positive recurrent state, all states are positive recurrent. In this case, we call such a Markov chain positive recurrent. One way to show the positive recurrence of a Markov chain is using the following Lyapunov-Foster criteria.

A function $L : \Omega \to \mathbb{R}_+$ is called a 'Lyapunov function' if $\sup_{\mathsf{x} \in \Omega} L(\mathsf{x}) = \infty$. The following theorem provides a criteria about the positive recurrence using a Lyapunov function, e.g., see [29].

THEOREM 1. *Consider a Lyapunov function $L : \Omega \to \mathbb{R}_+$ such that $B_\gamma := \{\mathsf{x} \in \Omega : L(\mathsf{x}) \leq \gamma\}$ is finite for any $\gamma \in \mathbb{R}_+$. Suppose there exist functions $\tau, \kappa : \Omega \to \mathbb{R}_+$ such that*

$$\mathbb{E}[L(\boldsymbol{X}(\tau(\mathsf{x}))) - L(\boldsymbol{X}(0)) \mid \boldsymbol{X}(0) = \mathsf{x}] \leq -\kappa(\mathsf{x}), \;\; \forall x \in \Omega,$$

and they satisfy the following conditions:

L1. $\inf_{\mathsf{x} \in \Omega} \kappa(\mathsf{x}) > -\infty$.

L2. $\liminf_{L(\mathsf{x}) \to \infty} \kappa(\mathsf{x}) > 0$.

L3. $\sup_{\mathsf{x} \in B_\gamma} \tau(\mathsf{x}) < \infty$ *for all* $\gamma \in \mathbb{R}_+$.

L4. $\limsup_{L(\mathsf{x}) \to \infty} \tau(\mathsf{x})/\kappa(\mathsf{x}) < \infty$.

Then, $\{\boldsymbol{X}(t)\}$ is positive recurrent.

3. MAIN RESULT

This section presents our main result, throughput-optimal and low-complexity scheduling algorithms. We first describe the generic concept of our algorithms utilizing certain oracle systems. The oracle systems are motivated by (randomized or deterministic) iterative algorithms which find an approximate optimal solution to the following optimization problem:

$$\max_{\boldsymbol{\rho} \in S} \boldsymbol{\rho} \cdot \boldsymbol{W}. \tag{1}$$

Typically, in an iterative algorithm, the output at each iteration depends only on information from the previous iteration (and weight $\boldsymbol{W}$). We call such information 'advice'.

One can understand an iteration of the algorithm as an interaction to an oracle system, where a query consisting of advice and weight is sent to an interactive oracle system, and the oracle returns a schedule together with an updated advice for the next iteration. This naturally suggests the following scheduling algorithm: at each time instance t, (a) it sends a query consisting of the advice $\boldsymbol{a}(t)$ produced in the previous step and weight $\boldsymbol{W}(t)$ decided by current queue lengths to the oracle system, and then (b) the oracle returns a schedule to use at the next time instance $t+1$. Hence, the complexity of the scheduling algorithm is that of the oracle system processing a query (equivalently, that of one iteration of an iterative algorithm), which is typically very small as we provide examples in Section 4. In the following sections, we provide the formal descriptions of oracle systems of interest and scheduling algorithms utilizing them.

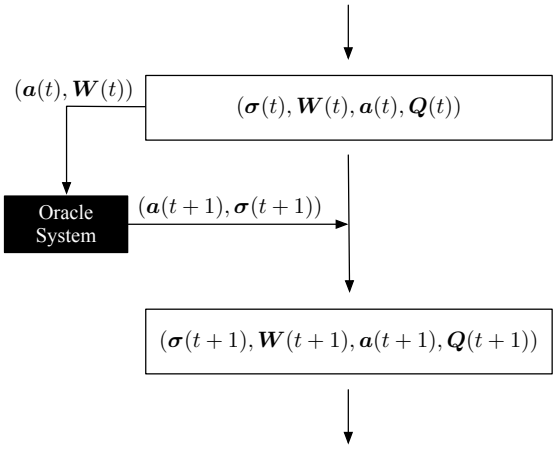

Figure 1: Scheduling with an interactive oracle system. At each time instance, a query consisting of current advice $\boldsymbol{a}(t)$ and weight $\boldsymbol{W}(t)$ is sent to the oracle. The oracle returns the schedule $\boldsymbol{\sigma}(t+1)$ to use at the next time instance and updated advice $\boldsymbol{a}(t+1)$.

3.1 Oracle System

We consider a network model introduced in Section 2.1 with set of queues V with $|V| = n$ and available schedules $S \subset \{0,1\}^n$. Then the oracle system we are interested in has an advice space $\mathbb{A}$ with the following input and output:

- It takes inputs $\boldsymbol{a} \in \mathbb{A}$ and $\boldsymbol{W} = [W_i \in \mathbb{Z}_+ : i \in V]$,
- It outputs (or returns) $\boldsymbol{\sigma} \in S$ and $\widehat{\boldsymbol{a}} \in \mathbb{A}$,

where $\boldsymbol{W}$, $\boldsymbol{a}$, and $\widehat{\boldsymbol{a}}$ are called 'weights' on queues, 'advice', and 'updated advice', respectively. Formally, let the outputs $\boldsymbol{\sigma} = \boldsymbol{\sigma}_{\mathsf{oracle}}(\boldsymbol{a}) = \boldsymbol{\sigma}_{\mathsf{oracle}}(\boldsymbol{W}, \boldsymbol{a})$ and $\widehat{\boldsymbol{a}} = \boldsymbol{a}_{\mathsf{oracle}}(\boldsymbol{a}) = \boldsymbol{a}_{\mathsf{oracle}}(\boldsymbol{W}, \boldsymbol{a})$, where the oracle can generate random outputs in general. We assume that the oracle system satisfies the following properties:

A1. For any $\eta, \delta \in (0,1)$, if $W_{\max} := \max_{i \in V} W_i$ is large enough, then there exists $h = h(W_{\max}, \eta, \delta)$ such that for any $t \geq h$ and advice $\boldsymbol{a} \in \mathbb{A}$,

$$\left(\boldsymbol{\sigma}_{\mathsf{oracle}}(\boldsymbol{a}_{\mathsf{oracle}}^{(t)}(\boldsymbol{a}))\right) \cdot \boldsymbol{W} \;\geq\; (1-\eta) \max_{\boldsymbol{\rho} \in S} \boldsymbol{\rho} \cdot \boldsymbol{W}, \quad \text{with probability} \geq 1-\delta$$

where $\boldsymbol{a}_{\mathsf{oracle}}^{(t)}$ is the function composing $\boldsymbol{a}_{\mathsf{oracle}}$ 't times', i.e., $\boldsymbol{a}_{\mathsf{oracle}}^{(t)} = \boldsymbol{a}_{\mathsf{oracle}}^{(t-1)} \circ \boldsymbol{a}_{\mathsf{oracle}}$.

A2. For any $\boldsymbol{W}$,

$$\left| \mathbb{A}(\boldsymbol{W}) := \bigcup_{\boldsymbol{a} \in \mathbb{A}} \boldsymbol{a}_{\mathsf{oracle}}(\boldsymbol{W}, \boldsymbol{a}) \right| \;<\; \infty. \tag{2}$$

The first assumption **A1** implies that the schedule $\boldsymbol{\sigma}$ generated by the oracle system after h interactions is an approximate solution to the optimization problem (1). The second assumption **A2** is merely for a technical purpose: it provides the countability of the underlying space of the network Markov chain which we define in the following section.

3.2 Scheduling Algorithm

In this section, we describe our scheduling algorithm given an oracle system. To this end, we let $f, g : \mathbb{R}_+ \to \mathbb{R}_+$ be positive-real-valued functions and

$$U_i(t) := \max\{f(Q_i(t)), g(Q_{\max}(t))\},$$

where $Q_{\max}(t) = \max_{i \in V} Q_i(t)$. The algorithm maintains parameters $\boldsymbol{W}(t) \in \mathbb{Z}_+^n$, $\boldsymbol{a}(t) \in \mathbb{A}$. It updates the schedule $\boldsymbol{\sigma}(t)$ and parameters $\boldsymbol{W}(t), \boldsymbol{a}(t)$ under an interaction with the oracle system as follows:

- $\boldsymbol{\sigma}(t+1) = \boldsymbol{\sigma}_{\mathsf{oracle}}(\boldsymbol{W}(t), \boldsymbol{a}(t))$,
- $\boldsymbol{a}(t+1) = \boldsymbol{a}_{\mathsf{oracle}}(\boldsymbol{W}(t), \boldsymbol{a}(t))$,
- $W_i(t+1)$ is the closest integer to $U_i(t+1)$ if
$$|W_i(t) - U_i(t+1)| \geq 2,$$
and $W_i(t+1) = W_i(t)$ otherwise.

Figure 1 illustrates the above scheduling algorithm using an interactive oracle system. For each time $t \in \mathbb{Z}_+$, we let

$$\boldsymbol{X}(t) := (\boldsymbol{\sigma}(t), \boldsymbol{W}(t), \boldsymbol{a}(t), \boldsymbol{Q}(t)) \in S \times \mathbb{Z}_+^n \times \mathbb{A} \times \mathbb{Z}_+^n.$$

Then, $\{\boldsymbol{X}(t) : t \in \mathbb{Z}_+\}$ becomes a Markov chain on the following state space Ω:

$$\Omega := \left\{ (\boldsymbol{\sigma}, \boldsymbol{W}, \boldsymbol{a}, \boldsymbol{Q}) : \begin{array}{l} U_i = \max\{f(Q_i), g(Q_{\max})\}, \\ |W_i - U_i| \leq 2, \\ \boldsymbol{a} \in \mathbb{A}(\boldsymbol{W}),\ \boldsymbol{\sigma} \in S, W_i, Q_i \in \mathbb{Z}_+, \end{array} \right\}.$$

Then, Ω is countable due to the assumption **A2** of the oracle system. Now we are ready to state the main result in this paper, where its proof is presented in Section 5.

THEOREM 2. *The above scheduling algorithm is throughput-optimal, i.e., $\{\boldsymbol{X}(t)\}$ is positive recurrent for any $\boldsymbol{\lambda} \in \Lambda^o$, if functions f, g, h satisfy the assumptions* **A1-A2** *and the following conditions:*

C1. *f and g are increasing, differentiable and concave.*

C2. $\lim_{x \to \infty} \frac{g(x)}{f(x)} = 0$ *and* $\lim_{x \to \infty} g(x) = \infty$.

C3. $f(0) = 0$.

C4. $\lim_{x \to \infty} f'(x) = \lim_{n \to \infty} g'(x) = 0$.

C5. *For any fixed* $\eta, \delta > 0$,

$$\lim_{x\to\infty} \frac{h(f(x),\eta,\delta)}{x} = 0.$$

C6. *There exists* $c \in (0,1)$ *such that, for any fixed* $\eta, \delta > 0$,

$$\lim_{x\to\infty} f'\left(f^{-1}\left(g\left((1-c)x\right)\right)\right) h\left(f((1+c)x),\eta,\delta\right) = 0.$$

In the following section, we provide several specific examples of throughput-optimal and low-complexity scheduling algorithms under Theorem 2.

4. EXAMPLES

In this section, we show the wide-applicability of our framework by designing several throughput-optimal and low-complexity scheduling algorithms under it. In particular, we provide examples of low-complexity oracle systems with corresponding explicit functions f, g, h so that they satisfy the assumptions **A1-A2** and the conditions **C1-C6** of Theorem 2.

4.1 Exhaustive Search (ES): Pick-and-Compare

The first oracle system we introduce utilizes an exhaustive search (ES) algorithm to find a maximum weight schedule $\boldsymbol{\sigma} \in S$, i.e., it maximizes $\boldsymbol{\sigma}\cdot\boldsymbol{W}$ over $\boldsymbol{\sigma} \in S$ for given V with $|V| = n$, $S \subset \{0,1\}^n$, and $\boldsymbol{W} \in \mathbb{Z}_+^n$. An ES algorithm keeps a search-index $\boldsymbol{\rho} \in \{0,1\}^n$ which may not be in S, and current maximum schedule $\boldsymbol{\sigma} \in S$. At each iteration, if $\boldsymbol{\rho}$ is in S and its weight is greater than the weight of $\boldsymbol{\sigma}$, replace $\boldsymbol{\sigma}$ by $\boldsymbol{\rho}$. In addition, it updates the search index $\boldsymbol{\rho}$ by $\boldsymbol{\rho}+1 \mod 2^n$ where $\boldsymbol{\rho} \in \{0,1\}^n$ can be considered as a non-negative integer less than 2^n and $\boldsymbol{\rho}+1 \mod 2^n$ is naturally defined under the interpretation.

ES oracle system. Formally, the ES oracle system has an advice space $\mathbb{A} = \{0,1\}^n \times S$. For an advice $\boldsymbol{a} = (\boldsymbol{\rho},\boldsymbol{\sigma}) \in \mathbb{A}$ and weights $\boldsymbol{W} \in \mathbb{Z}_+$, the oracle outputs $\boldsymbol{\sigma}_{\mathsf{oracle}}(\boldsymbol{W},\boldsymbol{a}) = \widehat{\boldsymbol{\sigma}}$ and $\boldsymbol{a}_{\mathsf{oracle}}(\boldsymbol{W},\boldsymbol{a}) = (\widehat{\boldsymbol{\rho}},\widehat{\boldsymbol{\sigma}})$ are decided as follows:

---

1. Set $\widehat{\boldsymbol{\sigma}} = \begin{cases} \boldsymbol{\rho} & \text{if } \boldsymbol{\rho}\in S \text{ and } \boldsymbol{\rho}\cdot\boldsymbol{W} > \boldsymbol{\sigma}\cdot\boldsymbol{W} \\ \boldsymbol{\sigma} & \text{otherwise} \end{cases}$,
2. Set $\widehat{\boldsymbol{\rho}} = \boldsymbol{\rho}+1 \mod 2^n$.

Since it compares all schedules in S exhaustively, after 2^n interactions with the oracle system, it returns a schedule $\boldsymbol{\sigma} \in S$ whose weight is maximum. Thus we have, for all $\boldsymbol{a} \in \mathbb{A}$ and $t \geq 2^n$,

$$\boldsymbol{\sigma}_{\mathsf{oracle}}\left(\boldsymbol{a}^{(t)}_{\mathsf{oracle}}(\boldsymbol{a})\right)\cdot\boldsymbol{W} \;\geq\; \max_{\boldsymbol{\rho}\in S}\boldsymbol{\rho}\cdot\boldsymbol{W},$$

which implies that we can take $h(W_{\max},\eta,\delta) = 2^n$ for the the assumption **A1**, i.e., h is a constant function with respect to $\boldsymbol{W},\eta,\delta$. In addition, the assumption **A2** holds obviously. Using these facts and Theorem 2, the following corollary can be derived.

COROLLARY 3. *The scheduling algorithm described in Section 3.2 using the ES oracle system is throughput-optimal if*

$$f(x) = x^a,\;\; g(x) = x^b,\;\; \textit{and } 0 < b < a < 1.$$

PROOF. It is elementary to check the conditions **C1-C5** of Theorem 2, where $h(W_{\max},\eta,\delta) = 2^n$. The condition **C6** of Theorem 2 can be derived as follows: for $0 < c < 1$,

$$\begin{aligned}&\lim_{x\to\infty} f'\left(f^{-1}\left(g\left((1-c)x\right)\right)\right) h\left(f((1+c)x),\eta,\delta\right)\\ &= \lim_{x\to\infty} 2^n\, a\, \left((1-c)\right)^{\frac{b(a-1)}{a}}\, x^{\frac{b(a-1)}{a}} \;=\; 0,\end{aligned}$$

This completes the proof of Corollary 3. □

4.2 Markov Chain Monte Carlo (MCMC)

The second oracle system we introduce is the Markov Chain Monte Carlo (MCMC) algorithm. For a detailed description of the oracle, we consider a special setup of available schedules S.

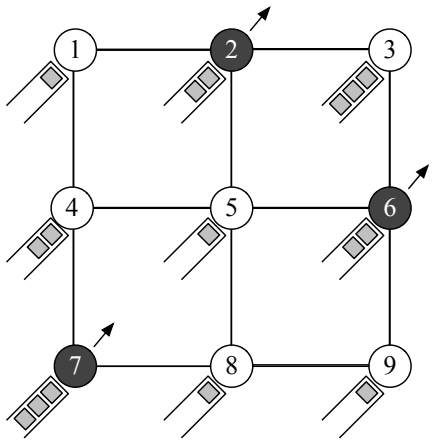

Figure 2: Wireless network with $n = 9$ queues in grid interference topology. Available schedules are $\{1,3,5,7,9\}$, $\{1,3,8\}$, $\{2,4,6,8\}$, $\{2,4,9\}$, $\{2,6,7\}$, $\{2,7,9\}$, etc.

Wireless network model. We consider an undirected graph $G = (V,\mathcal{E})$ and the set of available schedules S is defined as:

$$S \;=\; \{\boldsymbol{\sigma}\in\{0,1\}^n \;:\; \sigma_i+\sigma_j \leq 1,\;\; \forall\,(i,j)\in\mathcal{E}\}. \qquad (3)$$

The above scheduling constraint has been popularly used for modeling interference in wireless networks [10], i.e., two queues i and j cannot transmit simultaneously if they share an edge in the 'interference' graph G. We let $\mathcal{N}(i) = \{j : (i,j)\in\mathcal{E}\}$ denote the set of interfering neighbors of queue i.

MCMC oracle system. In the MCMC oracle system, the advice space is $\mathbb{A} = S$ and for an advice $\boldsymbol{a} = \boldsymbol{\sigma} \in \mathbb{A}$, the oracle outputs $\boldsymbol{\sigma}_{\mathsf{oracle}}(\boldsymbol{W},\boldsymbol{a}) = \widehat{\boldsymbol{\sigma}}$ and $\boldsymbol{a}_{\mathsf{oracle}}(\boldsymbol{W},\boldsymbol{a}) = \widehat{\boldsymbol{\sigma}}$ are decided randomly as follows:

---

1. Choose a queue i uniformly at random, and set
$$\widehat{\sigma}_j = \sigma_j, \qquad \text{for all } j \neq i.$$
2. If $\sigma_j = 0$ for some $j \in \mathcal{N}(i)$, then set $\widehat{\sigma}_i = 0$.
3. Otherwise, set
$$\widehat{\sigma}_i = \begin{cases} 1 & \text{with probability } \frac{\exp(W_i)}{1+\exp(W_i)} \\ 0 & \text{otherwise} \end{cases}.$$

Then, using known results on the mixing time of MCMC, one can prove that the assumption **A1** holds with

$$h(W_{\max}, \eta, \delta) = e^{C_1 W_{\max}} \left(C_2 + \log\left(\frac{1}{\eta\delta}\right)\right), \qquad (4)$$

where $C_1 = C_1(n), C_2 = C_2(n)$ are some ('n-dependent') constants independent of $W_{\max}$. The proof of (4) is a direct consequence of Lemma 3 and Lemma 7 in [12], and we omit the details due to space constraints. The assumption **A2** also holds obviously again. Using these facts and Theorem 2, the following corollary can be derived.

COROLLARY 4. *The scheduling algorithm described in Section 3.2 using the MCMC oracle system is throughput-optimal if*

$$f(x) = (\log(x+e))^a - 1 \quad \textit{and} \quad g(x) = (\log(x+e))^b - 1,$$

where $0 < a^2 < b < a < 1$.

PROOF. It is elementary to check the conditions **C1-C4** of Theorem 2. The condition **C5** is from (4) and $f(x) = (\log(x+e))^a - 1$:

$$\begin{aligned} &\frac{h(f(x),\eta,\delta)}{x} \\ &= \left(C_2 + \log\left(\frac{1}{\eta\delta}\right)\right)\left(e^{C_1(\log(x+e))^a - \log x - C_1}\right) \\ &\overset{x\to\infty}{\to} 0. \end{aligned}$$

Furthermore, the condition **C6** can be derived as follows:

$$\begin{aligned} &f'(f^{-1}(g((1-c)x)))\, h(f((1+c)x)) \\ &= \frac{a\,(C_2 + \log(1/(\eta\delta)))}{\left(\log((1-c)x+e)^{\frac{(1-a)b}{a}}\right)} \\ &\quad \times \left(e^{C_1(\log((1+c)x+e))^a - \log((1-c)x+e)^{\frac{b}{a}} - C_1}\right) \\ &\overset{x\to\infty}{\to} 0. \end{aligned}$$

This completes the proof of Corollary 4. □

We note that the scheduling algorithm described in Section 3.2 using the MCMC oracle system is the discrete-time version of the CSMA algorithm in [10, 12].

4.3 Belief Propagation (BP)

The third oracle system we introduce is the Belief Propagation (BP) algorithm, where BP is a popular heuristic iterative algorithm for solving inference problems arising in graphical models [30]. For the provable throughput-optimality of the scheduling algorithm using the BP oracle system, we consider a special setup of available schedules S which is described in what follows.

Input-queued switch model. We consider the undirected complete bipartite graph of left-vertices L, right-vertices R, and edges $\mathcal{E} = \{(i,j) : i \in L, j \in R\}$, where $|L| = |R| = m$ Then, the set of queues is $V = \mathcal{E}$ of $n = |V| \le m^2$ queues and

$$S = \left\{ \boldsymbol{\sigma} \in \{0,1\}^{\mathcal{E}} \;:\; \begin{array}{l} \sum_{j:(i,j)\in E} \sigma_{ij} \le 1 \;\forall i \in R, \\ \sum_{i:(i,j)\in E} \sigma_{ij} \le 1 \;\forall j \in L \end{array} \right\}, \qquad (5)$$

where this is a popular model for input-queued switches [21], and also a special case of the wireless network model described in the previous section.

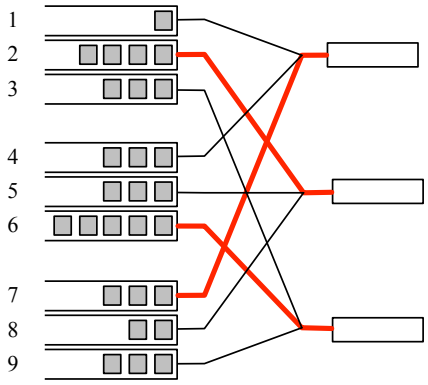

Figure 3: Input-queued switch with $m = 3$ input ports and $m = 3$ output ports, there are $m^2 = 9$ queues. Available schedules are $\{1,5,9\}$, $\{1,6,8\}$, $\{2,4,9\}$, $\{2,6,7\}$, $\{3,4,8\}$, and $\{3,5,7\}$.

BP oracle system. In the BP oracle system, the advice space is $\mathbb{A} = \mathbb{Z}_+^{|2\mathcal{E}|} \times S$ and for an advice $\boldsymbol{a} = (\boldsymbol{m}, \boldsymbol{\sigma}) \in \mathbb{A}$, where $\boldsymbol{m} = [m_{i\to j}, m_{j\to i} : (i,j) \in \mathcal{E}]$, the oracle outputs $\boldsymbol{\sigma}_{\mathsf{oracle}}(\boldsymbol{W}, \boldsymbol{a}) = \widehat{\boldsymbol{\sigma}}$ and $\boldsymbol{a}_{\mathsf{oracle}}(\boldsymbol{W}, \boldsymbol{a}) = (\widehat{\boldsymbol{m}}, \widehat{\boldsymbol{\sigma}})$ are decided as follows:

1. For each $(i,j) \in \mathcal{E}$, set

$$\begin{aligned} \widehat{\sigma}_{(i,j)} &= \begin{cases} 0 & \text{if } m_{i\to j} + m_{j\to i} > W'_{(i,j)} \\ 1 & \text{otherwise} \end{cases} \\ \widehat{m}_{i\to j} &= \max_{k\neq j:(i,k)\in\mathcal{E}} \left(W'_{(i,k)} - m_{k\to i}\right)_+, \end{aligned}$$

where

$$W'_{(i,j)} := W_{(i,j)} + r_{(i,j)} \quad (x)_+ := \begin{cases} x \text{ if } x \ge 0 \\ 0 \text{ otherwise} \end{cases}.$$

2. If $\widehat{\boldsymbol{\sigma}} \notin S$, reset $\widehat{\boldsymbol{\sigma}} = \boldsymbol{\sigma}$.

In the above, one can set $[r_{(i,j)}] \in [0,1]^{|\mathcal{E}|}$ such that $\boldsymbol{\sigma}^* \in \arg\max_{\boldsymbol{\sigma}} \boldsymbol{\sigma}\cdot\boldsymbol{W}'$ is unique, and $\xi \le \boldsymbol{\sigma}^*\cdot\boldsymbol{W}' - \max_{\boldsymbol{\sigma}\neq\boldsymbol{\sigma}^*} \boldsymbol{\sigma}\cdot\boldsymbol{W}'$ for some constant $\xi > 0$, e.g., let $r_{(i,j)}$ be a random number in the interval $[0,1]$. Then, from the works by Bayati et al. [31] or Sanghavi et al. [32], it follows that the assumption **A1** holds with

$$h = h(W_{\max}, \eta, \delta) = O(W_{\max}/\xi).$$

For the assumption **A2**, one can observe that $\mathbb{A}(\boldsymbol{W}) \subset \{0, 1, \ldots, W_{max}\}^{2|\mathcal{E}|} \times S$ which is obviously a finite set given $\boldsymbol{W}$. Using these facts and Theorem 2, the following corollary can be derived.

COROLLARY 5. *The scheduling algorithm described in Section 3.2 using the BP oracle system is throughput-optimal if*

$$f(x) = x^a,\ g(x) = x^b,\ \textit{and}\ 0 < \frac{a^2}{1-a} < b < a < \frac{1}{2}.$$

PROOF. It is elementary to check the conditions **C1-C5** of Theorem 2, where $h(W_{\max}, \eta, \delta) = O(W_{\max}/\xi)$. The condition **C6** of Theorem 2 can be derived as follows: for $c > 0$,

$$\begin{aligned}\lim_{x\to\infty} f'\left(f^{-1}\left(g\left((1-c)x\right)\right)\right) h\left(f((1+c)x), \eta, \delta\right)\\ = \lim_{x\to\infty} C \cdot x^{\frac{b(a-1)}{a}} \left((1+c)^a x^a + 1\right) = 0,\end{aligned}$$

where C is some constant depending on ξ, c, a, b and the last equality is from $0 < b, a < 1$ and $b > \frac{a^2}{1-a}$. This completes the proof of Corollary 5. □

We note that there are many different ways to design the BP oracle, for example, one can consider the following variant.

---

1. For each $(i,j) \in \mathcal{E}$, set
$$\begin{aligned}\widehat{\sigma}_{(i,j)} &= 0\\ b_{(i,j)} &= W'_{(i,j)} - m_{i\to j} - m_{j\to i}\\ \widehat{m}_{i\to j} &= \max_{k\neq j:(i,k)\in\mathcal{E}} \left(W'_{(i,k)} - m_{k\to i}\right)_+.\end{aligned}$$
2. Choose $(i,j) \in \mathcal{E}$ so that $b_{(i,j)}$ is the largest one among those which $\widehat{\boldsymbol{\sigma}} \in S$ after resetting $\widehat{\sigma}_{(i,j)} = 1$.
3. Reset $\widehat{\sigma}_{(i,j)} = 1$ and keep this procedure until no more edge is found in the step 2.

The scheduling algorithm described in Section 3.2 using the above second BP oracle system is also throughput-optimal for the same set of functions f, g in Theorem 5, where its proof is identical to that of the first BP oracle system. The main difference in the second BP oracle is that it decides the schedule $\widehat{\boldsymbol{\sigma}}(t)$ greedily based on the 'belief' $[b_{(i,j)}]$, while the first one simply checks whether $b_{(i,j)}$ is positive or not. A similar version of the scheduling algorithm using the second oracle system was actually studied in [33] heuristically, where by carefully choosing queue-length functions, Theorem 2 provides its formal throughput-optimality proof which is missing in [33].

4.4 Combination of BP and ES

In this section, we introduce the fourth oracle system combining two mechanisms, BP and ES, in the wireless interference network model described in Section 4.2. We call it the BP-ES oracle system.

BP-ES oracle system. In the oracle system, the advice space is $\mathbb{A} = \{0,1\}^n \times S \times \mathbb{R}^{2|\mathcal{E}|}$, and for an advice $\boldsymbol{a} = (\boldsymbol{\rho}, \boldsymbol{\sigma}, \boldsymbol{m}) \in \mathbb{A}$ where $\boldsymbol{m} = [m_{i\to j}, m_{i\to j} : (i,j) \in \mathcal{E}]$, the oracle outputs $\boldsymbol{\sigma}_{\mathsf{oracle}}(\boldsymbol{W}, \boldsymbol{a}) = \widehat{\boldsymbol{\sigma}}$ and $\boldsymbol{a}_{\mathsf{oracle}}(\boldsymbol{W}, \boldsymbol{a}) = (\widehat{\boldsymbol{\rho}}, \widehat{\boldsymbol{\sigma}}, \widehat{\boldsymbol{m}})$ are decided as follows:

---

1. For each edge $(i,j) \in \mathcal{E}$, set
$$\widehat{m}_{i\to j} = \left(W_i - \sum_{k\in\mathcal{N}(i)\setminus\{j\}} m_{k\to i}\right)_+.$$
2. Choose a queue i uniformly at random, and set
$$\widehat{\rho}_j = \rho_j, \qquad \text{for all } j \neq i, \qquad \widehat{\rho}_i = \rho_i + 1 \mod 2.$$
3. Set $\widehat{\boldsymbol{\sigma}} = \begin{cases}\boldsymbol{\rho} & \text{if } \boldsymbol{\rho} \in S \text{ and } \boldsymbol{\rho}\cdot\boldsymbol{W} > \max\{\boldsymbol{\sigma}\cdot\boldsymbol{W}, \boldsymbol{\delta}\cdot\boldsymbol{W}\}\\ \boldsymbol{\delta} & \text{else if } \boldsymbol{\delta} \in S \text{ and } \boldsymbol{\delta}\cdot\boldsymbol{W} > \boldsymbol{\sigma}\cdot\boldsymbol{W}\\ \boldsymbol{\sigma} & \text{otherwise}\end{cases}$,
where $\boldsymbol{\delta} = [\delta_i]$ is defined as
$$\delta_i = \begin{cases}1 & \text{if } W_i > \sum_{j\in\mathcal{N}(i)} m_{j\to i}\\ 0 & \text{otherwise}\end{cases}.$$

In the above, the second and third steps can be implementable in distributed fashions using Poisson clocks [10] and gossip algorithms for distributed averaging [34]. The BP-ES oracle system outputs schedule $\widehat{\boldsymbol{\sigma}}$ as the maximum weight one among three schedules $\boldsymbol{\sigma}$, $\boldsymbol{\rho}$ and $\boldsymbol{\delta}$, where $\boldsymbol{\sigma}$ is the current maximum schedule, $\boldsymbol{\rho}$ is the ES search index and $\boldsymbol{\delta}$ is decided by the BP mechanism. Since the search index $\boldsymbol{\rho}$ is decided randomly, the BP-ES oracle finds an maximum weight schedule with probability $\geq 1-\delta$ after $h(W_{\max}, \eta, \delta) = \log\delta/\log(1-2^{-n})$ interactions for the the assumption **A1**. The assumption **A2** holds with similar reasons as before. Using these facts and Theorem 2, the following corollary can be derived.

COROLLARY 6. *The scheduling algorithm described in Section 3.2 with the BP-ES oracle system is throughput-optimal if*

$$f(x) = x^a, \qquad g(x) = x^b, \qquad \textit{and} \qquad 0 < b < a < 1.$$

PROOF. The proof of Corollary 6 is almost identical to that of Corollary 3. □

5. PROOF OF MAIN THEOREM

This section provides the proof of Theorem 2, i.e., we show that $\{\boldsymbol{X}(t) : t \in \mathbb{Z}_+\}$ described in Section 3.2 is positive recurrent, as long as $\boldsymbol{\lambda} \in \Lambda^o$ and f, g, h satisfy the assumptions **A1-A2** and the conditions **C1-C6** stated in Theorem 2. Towards this, we shall construct a Lyapunov function L and auxiliary functions τ, κ so that they satisfy the conditions of Theorem 1.

First, note that the state space Ω is countably infinite from the assumption **A2** of the oracle system. For $\mathsf{x} = (\boldsymbol{\sigma}, \boldsymbol{W}, \boldsymbol{a}, \boldsymbol{Q}) \in \Omega$, we consider the following function:

$$L(\mathsf{x}) := \sum_{i=1}^n F(Q_i),$$

where $F(x) = \int_0^x f(s)ds$. Since $\lim_{x\to\infty} f(x) = \infty$ from the condition **C2**, we have $L(\mathsf{x}) \to \infty$ as $Q_{\max} = \max_{i\in V} Q_i \to \infty$. Thus, $\sup_{\mathsf{x}\in\Omega} L(\mathsf{x}) = \infty$ and $L : \Omega \to \mathbb{R}_+$ is a Lyapunov function. With an abuse of notation, we use $L(t)$ to denote $L(\boldsymbol{X}(t))$.

Now, given $\boldsymbol{\lambda} \in \Lambda^o$, there exists $\varepsilon > 0$ and $[\alpha_{\boldsymbol{\rho}} : \boldsymbol{\rho} \in S] \in [0,1]^{|S|}$ so that

$$\sum_{\boldsymbol{\rho}\in S} \alpha_{\boldsymbol{\rho}} = 1 - \varepsilon < 1 \quad \text{and} \quad \boldsymbol{\lambda} < \sum_{\boldsymbol{\rho}\in S} \alpha_{\boldsymbol{\rho}}\boldsymbol{\rho}. \tag{6}$$

For a state $\mathsf{x} = (\boldsymbol{\sigma}, \boldsymbol{W}, \boldsymbol{a}, \boldsymbol{Q}) \in \Omega$, we consider the following functions τ and κ:

$$\tau(\mathsf{x}) = \min\left\{ \left\lfloor \frac{1}{f'\left(f^{-1}((1-c)Q_{\max})\right)} \right\rfloor, \lfloor c\, Q_{\max} \rfloor \right\},$$
$$\frac{\kappa(\mathsf{x})}{\tau(\mathsf{x})} = \left(\frac{\varepsilon}{2}(1-\beta)(1-\alpha) - \frac{2n(\beta + (1-\beta)\alpha)}{1-c} \right) \times f((1-c)Q_{\max}) - n\, f'(0),$$

where c is a constant in the condition **C6**, $\lfloor x \rfloor$ is the largest integer not greater than x, and $\alpha, \beta \in (0,1)$ are some constants satisfying

$$\frac{\varepsilon}{2}(1-\beta)(1-\alpha) - \frac{2n(\beta + (1-\beta)\alpha)}{1-c} > 0.$$

For example, one can choose $\alpha = \beta = \frac{\varepsilon(1-c)}{32n}$. Using the above functions, we shall establish the following lemma.

LEMMA 7. *For every initial state $\mathsf{x} \in \Omega$ with large enough $L(\mathsf{x})$ and every arrival rate $\boldsymbol{\lambda} \in \Lambda^o$, it follows that*

$$\mathbb{E}[L(\tau(\mathsf{x})) - L(0)] \,|\, \boldsymbol{X}(0) = \mathsf{x}] \le -\kappa(\mathsf{x}). \tag{7}$$

The proof of the above lemma is presented in Section 5.1. In the remainder of this section, we check that τ and κ satisfy the desired conditions **L1-L4** of Theorem 1, which leads to the positive recurrence of the Markov chain $\{\boldsymbol{X}(t) : t \in \mathbb{Z}_+\}$.

First, it is easy to check that $B_\gamma = \{\mathsf{x} \in \Omega : L(\mathsf{x}) \le \gamma\}$ is finite. Without loss of generality, we can also assume that (7) holds for every $\mathsf{x} \in \Omega$ by letting $\tau(\mathsf{x}) = \kappa(\mathsf{x}) = 0$ if $L(\mathsf{x})$ is not large enough. From the conditions **C2** and **C4**, we have $\tau(\mathsf{x}) \to \infty$ as $Q_{\max} \to \infty$. Since f is an unbounded increasing function due to **C1** and **C2**, we have $\kappa(\mathsf{x}) \to \infty$ as $Q_{\max} \to \infty$. Thus, the first two conditions **L1** and **L2** are satisfied.

By the definition of function τ, it is not hard to check that τ increases as $Q_{\max}$ does. Therefore, it follows that

$$\sup_{L(\mathsf{x}) \le \gamma} \tau(\mathsf{x}) \le \sup_{Q_{\max} \le F^{-1}(\gamma)} \tau(\mathsf{x}) = \tau(F^{-1}(\gamma)) < \infty,$$

which implies the condition **L3**. Finally, as $L(\mathsf{x}) \to \infty$ (i.e., $Q_{\max} \to \infty$),

$$\frac{\tau(\mathsf{x})}{\kappa(\mathsf{x})} = \left[\left(\frac{\varepsilon}{2}(1-\beta)(1-\alpha) - \frac{2n(\beta + (1-\beta)\alpha)}{1-c} \right) \times f((1-c)Q_{\max})\tau(\mathsf{x}) - n\, f'(0) \right]^{-1} \to 0,$$

which implies the condition **L4**. This completes the proof of Theorem 2.

5.1 Proof of Lemma 7

To begin with, we describe the main steps to prove Lemma 7 at a high level:

1. For given initial state x with large enough $Q_{\max}$, $\boldsymbol{W}(t)$ does not change many times in $[0, \tau(\mathsf{x})]$, which is formally stated in Lemma 8.

2. The schedule $\boldsymbol{\sigma}(t)$ is close to a max-weight one with respect to $f(\boldsymbol{Q}(t))$ 'mostly' in the time interval $[0, \tau(\mathsf{x}))]$, which is formally stated in Lemma 9.

The first step plays a key role for establishing the second step: since $\boldsymbol{W}(t)$ does not change often, the interactive oracle system can find a max-weight schedule with respect to $\boldsymbol{W}(t) \approx f(\boldsymbol{Q}(t))$. The max-weight property in the second step will lead to the negative drift property in Lemma 7.

Now we proceed toward proving Lemma 7 formally. Let T_m is the time at which the m-th change of $\boldsymbol{W}(t)$ occurs, i.e., $\boldsymbol{W}(t)$ remains fixed during the time interval $[T_m, T_{m+1})$. Formally, let $T_0 = 0$ and for $m \ge 1$, iteratively define

$$T_m := \inf\{t \in \mathbb{Z}_+ : \boldsymbol{W}(t-1) \ne \boldsymbol{W}(t),\ t > T_{m-1}\}.$$

In the following lemma, for initial state x with large enough $Q_{\max}$, we show that $\boldsymbol{W}(t)$ can be changed at most n times in time interval $[0, \tau(\mathsf{x})]$, i.e., $T_{n+1} > \tau(\mathsf{x})$.

LEMMA 8. *For any initial state $\mathsf{x} = (\boldsymbol{\sigma}, \boldsymbol{W}, \boldsymbol{a}, \boldsymbol{Q}) \in \Omega$ with large enough $Q_{\max}$, the following event occurs with probability 1: $T_{n+1} > \tau(\mathsf{x})$ and*

$$(1-c) \ \le \ \frac{Q_{\max}(t)}{Q_{\max}} \ \le \ (1+c), \qquad \text{for all } t \in [0, \tau(\mathsf{x})], \tag{8}$$

*where c is the constant in the condition **C6** of Theorem 2.*

Since $\boldsymbol{W}(t)$ remains fixed for $t \in [T_m, T_{m+1})$, the assumption **A1** implies that that with high probability, $\boldsymbol{\sigma}(t)$ is close to the max-weight schedule with respect to $\boldsymbol{W}(t)$ for $T_m + h \le t \le T_{m+1}$. Using this observation with Lemma 8, we obtain the following lemma that with high probability, the schedule $\boldsymbol{\sigma}(t)$ is closed to a max-weight one with respect to $f(\boldsymbol{Q}(t))$ 'mostly' in the time interval $[0, \tau(\mathsf{x}))]$.

LEMMA 9. *For any $\eta, \alpha, \beta \in (0,1)$ and initial state $\mathsf{x} = (\boldsymbol{\sigma}, \boldsymbol{W}, \boldsymbol{a}, \boldsymbol{Q}) \in \Omega$ with large enough $Q_{\max}$, it follows that*

$$\mathbb{P}\big[\, |T(\mathsf{x}, \eta)| \ge (1-\alpha)\ \tau(\mathsf{x}) \ \mid \ \boldsymbol{X}(0) = \mathsf{x}\big] \ \ge \ 1 - \beta,$$

where

$$T(\mathsf{x}, \eta) := \Big\{ t \in [0, \tau(\mathsf{x})] \ : \ \boldsymbol{\sigma}(t) \cdot f(\boldsymbol{Q}(t)) \ge (1-\eta) \max_{\boldsymbol{\rho} \in S} \boldsymbol{\rho} \cdot f(\boldsymbol{Q}(t)) \Big\}.$$

In the remainder of this section, we shall establish the proof of Lemma 7 using Lemma 8 and Lemma 9. To this end, we first recall that

$$Q_{i+1}(t) - Q_i(t) = A_i(t) - \sigma_i(t) I_{\{Q_i(t) > 0\}}.$$

From the conditions **C1** and **C4**, we have

$$\begin{aligned} & L(t+1) - L(t) \\ &= \sum_{i=1}^n \int_{Q_i(t)}^{Q_i(t+1)} f(s) ds \\ &\le \sum_{i=1}^n \big(f(Q_i(t)) + f'(0) \big) \left(Q_i(t+1) - Q_i(t) \right) \\ &\le \boldsymbol{A}(t) \cdot f(\boldsymbol{Q}(t)) - \boldsymbol{\sigma}(t) \cdot f(\boldsymbol{Q}(t)) + n f'(0), \end{aligned} \tag{9}$$

where $\boldsymbol{A}(t) = [A_i(t) : i \in V]$. In addition, from (6), we have an upper bound on the expectation of the first term in (9):

$$\begin{aligned} & \mathbb{E}[\boldsymbol{A}(t) \cdot f(\boldsymbol{Q}(t)) \mid \boldsymbol{X}(0) = \mathsf{x}] \\ &= \mathbb{E}[\boldsymbol{A}(t)] \cdot \mathbb{E}[f(\boldsymbol{Q}(t)) \mid \boldsymbol{X}(0) = \mathsf{x}] \\ &\le \sum_{\boldsymbol{\rho} \in S} \alpha_{\boldsymbol{\rho}} \boldsymbol{\rho} \cdot \mathbb{E}[f(\boldsymbol{Q}(t)) \mid \boldsymbol{X}(0) = \mathsf{x}] \\ &\le (1-\varepsilon) \mathbb{E}[\max_{\boldsymbol{\rho} \in S} \boldsymbol{\rho} \cdot f(\boldsymbol{Q}(t)) \mid \boldsymbol{X}(0) = \mathsf{x}]. \end{aligned} \tag{10}$$

By combining (9) and (10), it follows that

$$\begin{aligned}
&\mathbb{E}[L(\tau(\mathsf{x})) - L(0)] \,|\, \boldsymbol{X}(0) = \mathsf{x}] \\
&= \mathbb{E}\left[\sum_{t=0}^{\tau(\mathsf{x})-1} L(t+1) - L(t) \,\middle|\, \boldsymbol{X}(0) = \mathsf{x}\right] \\
&\le \sum_{t=0}^{\tau(\mathsf{x})-1} \mathbb{E}\left[\boldsymbol{A}(t) \cdot f(\boldsymbol{Q}(t)) \,|\, \boldsymbol{X}(0) = \mathsf{x}\right] \\
&\quad - \sum_{t=0}^{\tau(\mathsf{x})-1} \mathbb{E}\left[\boldsymbol{\sigma}(t) \cdot f(\boldsymbol{Q}(t)) \,|\, \boldsymbol{X}(0) = \mathsf{x}\right] + n f'(0)\tau(\mathsf{x}) \\
&= \mathbb{E}\left[\sum_{t=0}^{\tau(\mathsf{x})-1} \Delta(t) \,\middle|\, \boldsymbol{X}(0) = \mathsf{x}\right],
\end{aligned}$$

where

$$\Delta(t) := (1-\varepsilon) \max_{\boldsymbol{\rho} \in S} \boldsymbol{\rho} \cdot f(\boldsymbol{Q}(t) - \boldsymbol{\sigma}(t) \cdot f(\boldsymbol{Q}(t)).$$

Now we will bound $\Delta(t)$. From Lemma 9 with $\eta = \varepsilon/2$, where ε is defined in (6), it follows that for $t \in T(\mathsf{x}, \eta)$,

$$\begin{aligned}
\Delta(t) &= (1-\varepsilon) \max_{\boldsymbol{\rho} \in S} \boldsymbol{\rho} \cdot f(\boldsymbol{Q}(t) - \boldsymbol{\sigma}(t) \cdot f(\boldsymbol{Q}(t)) \\
&\le -\frac{\varepsilon}{2} \max_{\boldsymbol{\rho} \in S} \boldsymbol{\rho} \cdot f(\boldsymbol{Q}(t)) \\
&\le \frac{\varepsilon}{2} f((1-c)Q_{\max}), \qquad (11)
\end{aligned}$$

where the last inequality is from Lemma 8. On the other hand, for any $t \in [0, \tau(\mathsf{x})]$, we know that

$$\Delta(t) \;\le\; n\, f((1+c)Q_{\max}).$$

Using these facts with Lemma 9, we further have that

$$\begin{aligned}
&\mathbb{E}[L(\tau(\mathsf{x})) - L(0)] \,|\, \boldsymbol{X}(0) = \mathsf{x}] \\
&\le \mathbb{E}\left[\sum_{t=0}^{\tau(\mathsf{x})-1} \Delta(t) \,\middle|\, \boldsymbol{X}(0) = \mathsf{x}\right] \\
&= \Pr[\mathcal{A}^c] \cdot \mathbb{E}\left[\sum_{t=0}^{\tau(\mathsf{x})-1} \Delta(t) \,\middle|\, \boldsymbol{X}(0) = \mathsf{x}, \mathcal{A}^c\right] \\
&\qquad + \Pr[\mathcal{A}] \cdot \mathbb{E}\left[\sum_{t=0}^{\tau(\mathsf{x})-1} \Delta(t) \,\middle|\, \boldsymbol{X}(0) = \mathsf{x}, \mathcal{A}\right] \\
&\le \beta\, n\, f((1+c)Q_{\max})\tau(\mathsf{x}) + n\, f'(0)\tau(\mathsf{x}) \\
&\qquad + (1-\beta)\, n\, f((1+c)Q_{\max})\alpha\, \tau(\mathsf{x}) \\
&\qquad + (1-\beta)\left(-\frac{\varepsilon}{2} f((1-c)Q_{\max})\right)(1-\alpha)\tau(\mathsf{x}) \\
&\le \left(-\frac{\varepsilon}{2}(1-\beta)(1-\alpha) + \frac{2n(\beta + (1-\beta)\alpha)}{1-c}\right) \times \\
&\qquad f((1-c)Q_{\max})\tau(\mathsf{x}) + n\, f'(0)\tau(\mathsf{x}) \\
&= -\kappa(\mathsf{x}),
\end{aligned}$$

where $\mathcal{A}$ denotes the event that $|T(\mathsf{x}, \eta)| \ge (1-\alpha)\tau(\mathsf{x})$ and in the last inequality, we use the following property for concave function f with $f(0) = 0$:

$$\frac{f((1+c)x)}{f((1-c)x)} \le \frac{2}{1-c} \frac{f((1+c)x)}{f(2x)} \le \frac{2}{1-c}.$$

This completes the proof of Lemma 7.

5.2 Proof of Lemma 8

We first proceed toward proving (8) in Lemma 8. It follows from the following observations:

$$\begin{aligned}
|Q_i(t+1) - Q_i(t)| \;&\le\; 1, \\
\tau(\mathsf{x}) \;&\le\; c\, Q_{\max},
\end{aligned}$$

where the first inequality is from our modeling assumption that $A_i(t)$, $\sigma_i(t) \in \{0, 1\}$ and $Q_i(t+1) - Q_i(t) = A_i(t) - \sigma_i(t) I_{\{Q_i(t) \ge 1\}}$, and the second inequality is from the definition of τ. Hence, for $t \in [0, \tau(\mathsf{x})]$, $Q_{\max}(t)$ does not change more than $c\, Q_{\max}$ from its initial value $Q_{\max}(0) = Q_{\max}$, which completes the proof of (8).

Now, we aim for proving $T_{n+1} > \tau(\mathsf{x})$. It suffices to show that, for given initial state x with large enough $Q_{\max}$,

$$|U_i(t) - U_i(0)| \;\le\; 1, \quad \forall i \in V \text{ and } \forall t \in [0, \tau(\mathsf{x})], \qquad (12)$$

i.e., $U_i(t)$ does not change more than 1, and hence $W_i(t)$ can be updated at most once for each $i \in V$. To prove (12), it is enough to show the following inequality:

$$|U_i(t+1) - U_i(t)| \;\le\; f'(f^{-1}(g((1-c)\, Q_{\max}))), \qquad (13)$$

since using the definition of τ,

$$\begin{aligned}
|U_i(t) - U_i(0)| \;&\le\; f'(f^{-1}(g((1-c)\, Q_{\max}))) \cdot t \\
&\le\; f'(f^{-1}(g((1-c)\, Q_{\max}))) \cdot \tau(\mathsf{x}) \\
&\le\; 1.
\end{aligned}$$

To this end, for $t \in [0, \tau(\mathsf{x})]$, consider following cases:

1. If $f(Q_i(t)) > g(Q_{\max}(t))$, then from the definition

   $$U_i(t) = \max\{f(Q_i(t)), g(Q_{\max}(t))\},$$

   the amount of change $|U_i(t+1) - U_i(t)|$ is bounded by

   $$\begin{aligned}
   f'(Q_i(t)) &\le f'(f^{-1}(g(Q_{\max}(t)))) \\
   &\le f'(f^{-1}(g((1-c)\, Q_{\max}))),
   \end{aligned}$$

   where we use (8) and the condition **C1**, i.e., f' is decreasing.

2. Otherwise, it is bounded by

   $$\begin{aligned}
   g'(Q_{\max}(t)) &\le f'(Q_{\max}(t)) \\
   &\le f'((1-c)\, Q_{\max}) \\
   &\le f'(f^{-1}(g((1-c)\, Q_{\max}))),
   \end{aligned}$$

   where we use (8) and the conditions **C1-C2**, i.e., f' is decreasing, $f(x) > g(x)$ and $f'(x) > g'(x)$ for large enough x. This completes the proof of (13) and hence that of Lemma 8.

5.3 Proof of Lemma 9

Without loss of generality, assume that $\eta \le \frac{7}{8}$ and let

$$\begin{aligned}
\eta' &= 1 - (1-\eta)^{1/3} \le \frac{1}{2}, \qquad & \alpha' &= 1 - \sqrt{1-\alpha}, \\
\gamma &= \frac{\beta}{n+1}, & \delta &= \alpha'\gamma.
\end{aligned}$$

To simplify notation, we will use $h(x)$ instead of $h(x, \eta', \delta)$. From conditions **C1-C6**, it is elementary to check that for

large enough x,

$$\frac{2n}{f\left((1-c)x\right)} \le \eta', \quad (14)$$

$$\frac{n\, g((1+c)x)+2n}{\frac{1}{2}(f((1-c)x)-2)} \le \eta', \quad (15)$$

$$\frac{(n+1)\, h(f((1+c)x+2))}{cx} \le \alpha', \quad (16)$$

$$(n+1)\, f'(f^{-1}(g((1-c)x)))\, h(f((1+c)x+2)) \le \alpha', \quad (17)$$

where their detailed proofs are given in Appendix A. Due to Lemma 8, we can assume that $T_{n+1} > \tau(\mathsf{x})$ and

$$(1-c)Q_{\max} \;\le\; Q_{\max}(t) \;\le\; (1+c)Q_{\max}. \quad (18)$$

The sufficient conditions to prove Lemma 9 are the following:

(a) For all $t \in [0, \tau(\mathsf{x})]$,

$$(1-\eta')\left[\max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot f(\boldsymbol{Q}(t))\right] \le \max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot \boldsymbol{W}(t).$$

(b) With probability $\ge 1-\beta$, at least $(1-\alpha)\tau(\mathsf{x})$ number of time instance $t \in [0, \tau(\mathsf{x})]$ satisfy

$$(1-\eta')\left[\max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot \boldsymbol{W}(t)\right] \le \boldsymbol{\sigma}(t)\cdot \boldsymbol{W}(t). \quad (19)$$

(c) For all $t \in [0, \tau(\mathsf{x})]$ at which (19) is satisfied,

$$(1-\eta')\left(\boldsymbol{\sigma}(t)\cdot \boldsymbol{W}(t)\right) \le \boldsymbol{\sigma}(t)\cdot f(\boldsymbol{Q}(t)).$$

The proof of Lemma 9 comes immediately from (a), (b) and (c) as follows:

$$\begin{aligned}(1-\eta)\left[\max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot f(\boldsymbol{Q}(t))\right] &= (1-\eta')^3\left[\max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot f(\boldsymbol{Q}(t))\right]\\ &\le (1-\eta')^2\left[\max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot \boldsymbol{W}(t)\right]\\ &\le (1-\eta')\left(\boldsymbol{\sigma}(t)\cdot \boldsymbol{W}(t)\right)\\ &\le \boldsymbol{\sigma}(t)\cdot f(\boldsymbol{Q}(t)),\end{aligned}$$

where with probability $\ge 1-\beta$, at least $(1-\alpha)\tau(\mathsf{x})$ number of time instance $t \in [0, \tau(\mathsf{x})]$ satisfy the second last and last inequalities. Hence, we proceed toward proving (a), (b) and (c).

Proof of (a). Recall that our scheduling algorithm in Section 3.2 maintains $|U_i(t) - W_i(t)| \le 2$, where $U_i(t) = \max\{f(Q_i(t)), g(Q_{\max}(t))\}$. Thus, for $t \in \mathbb{Z}_+$ and $i \in V$, we have

$$f(Q_i(t)) - 2 \;\le\; W_i(t) \;\le\; f(Q_i(t)) + g(Q_{\max}(t)) + 2, \quad (20)$$

and

$$\max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot f(\boldsymbol{Q}(t)) - 2n \;\le\; \max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot \boldsymbol{W}(t).$$

For large enough $Q_{\max}$, combining above inequality with (14), we conclude that

$$\begin{aligned}&(1-\eta')\left[\max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot f(\boldsymbol{Q}(t))\right]\\ &\le \left(1-\frac{2n}{f((1-c)Q_{\max})}\right)\left[\max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot f(\boldsymbol{Q}(t))\right]\\ &\le \max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot f(\boldsymbol{Q}(t)) - 2n\\ &\le \max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot \boldsymbol{W}(t),\end{aligned}$$

where we use $f((1-c)Q_{\max}) \le \max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot f(\boldsymbol{Q}(t))$ for $t \in [0, \tau(\mathsf{x})]$ from (8).

Proof of (b). We first state the following key proposition for the proof of (b).

PROPOSITION 10. *Suppose the assumption* **A1** *holds. Then, for any advice* $\boldsymbol{a} \in \mathbb{A}$ *and* $l \ge h = h(W_{\max}, \eta, \delta)$*, the following event occurs with probability* $\ge 1-\gamma$*:*

$$\left|\left\{t \in [h, l) \;:\; \left(\boldsymbol{\sigma}_{\mathsf{oracle}}(\boldsymbol{a}^{(t)}_{\mathsf{oracle}}(\boldsymbol{a}))\right)\cdot \boldsymbol{W} \;\ge\; (1-\eta)\max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot \boldsymbol{W}\right\}\right| \;\ge\; (1-\delta/\gamma)(l-h)$$

PROOF. For $t \ge h$, define a binary random variable $Z_t \in \{0,1\}$ such that $Z_t = 1$ if

$$\left(\boldsymbol{\sigma}_{\mathsf{oracle}}(\boldsymbol{a}^{(t)}_{\mathsf{oracle}}(\boldsymbol{a}))\right)\cdot \boldsymbol{W} \;<\; (1-\eta)\max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot \boldsymbol{W},$$

and $Z_t = 0$ otherwise. Then, from the assumption **A1**, $\mathbb{E}[Z_t] < \delta$ and $\mathbb{E}\left[\sum_{t=h}^{l-1} Z_t\right] < \delta(l-h)$. The conclusion of Proposition 10 directly follows from the Markov inequality applying to the random variable $\sum_{t=h}^{l-1} Z_t$. □

From the above proposition, it follows that with probability $\ge 1-\gamma$,

$$\left|\left\{t \in [T_i + h, T_{i+1}) : \boldsymbol{\sigma}(t)\cdot \boldsymbol{W}(t) \ge (1-\eta')\max_{\boldsymbol{\rho}\in S} \boldsymbol{\rho}\cdot \boldsymbol{W}(t)\right\}\right| \ge (1-\alpha')(T_{i+1} - T_i - h),$$

where from Lemma 8, we set

$$h \;=\; h((1+c)Q_{\max}+2) \;=\; h((1+c)Q_{\max}+2, \eta, \delta),$$

due to $W_{\max}(t) \le (1+c)Q_{\max} + 2$. Since $\boldsymbol{W}(t)$ changes at most n times in $[0, \tau(\mathsf{x})]$ from Lemma 8, one can use the union bound and conclude that with probability $\ge 1-\beta = 1-(n+1)\gamma$,

$$(1-\eta')\max \boldsymbol{\rho}\cdot \boldsymbol{W}(t) \le \boldsymbol{\sigma}(t)\cdot \boldsymbol{W}(t),$$

for at least $(1-\alpha')$ fraction of times in $\bigcup_{i=0}^{n}[T_i + h, T_{i+1})$. Furthermore, from (16), (17), and the definition of τ, we have

$$\begin{aligned}(n+1)h &= (n+1)h(f(1-c)Q_{\max}+2)\\ &\le \frac{\alpha'}{2}\min\left\{c\, Q_{\max}, \frac{1}{f'(f^{-1}(g(1-c)Q_{\max}))}\right\}\\ &\le \frac{\alpha'}{2}(\tau(\mathsf{x})+1) \;\le\; \alpha'\tau(\mathsf{x}).\end{aligned}$$

Thus, it follows that

$$\begin{aligned}\left|\bigcup_{i=0}^{n}[T_i + h, T_{i+1})\right| &\ge \tau(\mathsf{x}) - (n+1)h\\ &\ge (1-\alpha')\tau(\mathsf{x}).\end{aligned}$$

Therefore, with probability $\ge 1-\beta$, for at least $(1-\alpha')^2 = (1-\alpha)$ fraction of times in the interval $[0, \tau(\mathsf{x})]$, (19) holds.

Proof of (c). From (20), we first have that

$$\boldsymbol{\sigma}(t)\cdot \boldsymbol{W}(t) \;\le\; \boldsymbol{\sigma}(t)\cdot f(\boldsymbol{Q}(t)) + n\, g((1+c)Q_{\max}) + 2n.$$

For all $t \in [0, \tau(\mathsf{x})]$ at which (19) is satisfied,

$$\begin{aligned}\frac{1}{2}\left(f((1-c)Q_{\max})-2\right) &\leq (1-\eta')\left[\max_{\boldsymbol{\rho}\in S}\boldsymbol{\rho}\cdot\boldsymbol{W}(t)\right]\\ &\leq \boldsymbol{\sigma}(t)\cdot\boldsymbol{W}(t),\end{aligned}$$

where we use (8), (20) and $\eta' \leq 1/2$. For large enough $Q_{\max}$, combining the above inequality with (15) leads to (c):

$$\begin{aligned}&(1-\eta')\boldsymbol{\sigma}(t)\cdot\boldsymbol{W}(t)\\ &\leq \left(1-\frac{n\,g((1+c)Q_{\max})+2n}{\frac{1}{2}(f((1-c)Q_{\max})-2)}\right)\boldsymbol{\sigma}(t)\cdot\boldsymbol{W}(t)\\ &\leq \boldsymbol{\sigma}(t)\cdot\boldsymbol{W}(t)-n\,g((1+c)Q_{\max})+2n\\ &\leq \boldsymbol{\sigma}(t)\cdot f(\boldsymbol{Q}(t)).\end{aligned}$$

This completes the proof of Lemma 9.

6. EXPERIMENTAL RESULTS

This section presents experimental results for the performance of our scheduling algorithm using various interactive oracle systems. We consider two network models: input-queued switch networks and wireless networks described in Section 4.2 and Section 4.3, respectively. We compare the average queue-length in our scheduling algorithm using the MCMC, ES and BP oracle systems with that of the MW algorithm [1].

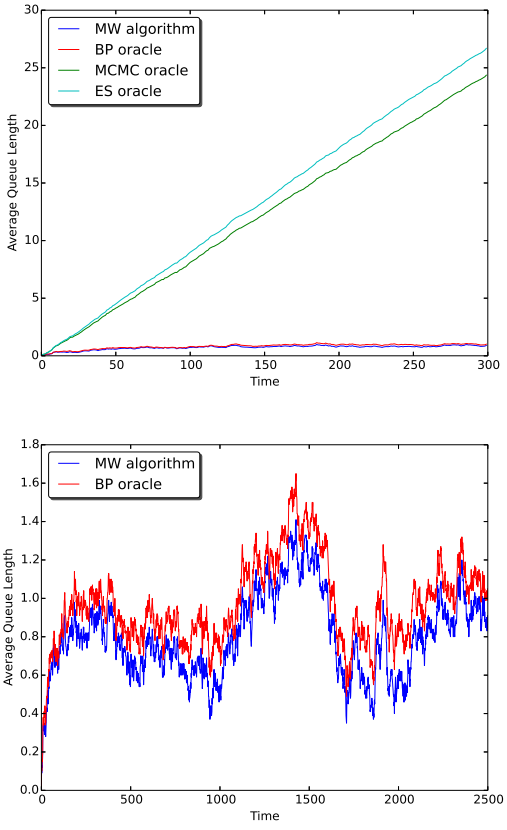

Figure 4: Performances of scheduling algorithms using MCMC, BP and ES oracle systems in input-queued switch network.

Input-queued switch network. We first consider an input-queued switch network on $m \times m$ complete bipartite graph topology in Section 4.3. We report the average queue lengths in the switch network for $m = 10$ with load 90%, i.e., arrival rates $\lambda_{(i,j)} = 0.9/m$ for all $(i,j) \in \mathcal{E}$ and any algorithm cannot stabilize the queues if $\lambda_{(i,j)} > 1/m$ for all $(i,j) \in \mathcal{E}$. We run the MW algorithm and our low-complexity scheduling algorithms using the MCMC, ES, and BP oracle systems, where we consider the second BP oracle described in Section 4.3. Here we use $f(x) = x^{1/4}, g(x) = x^{1/10}$ for ES and BP oracles and $f(x) = \log(x+1), g = \log(x+1)^{1/10}$ for MCMC oracle. Figure 4 summarizes our experimental results and we observe that the performance of our scheduling algorithm using BP oracle system is as good as that of the MW algorithm, although it is much simpler to code and faster to run.

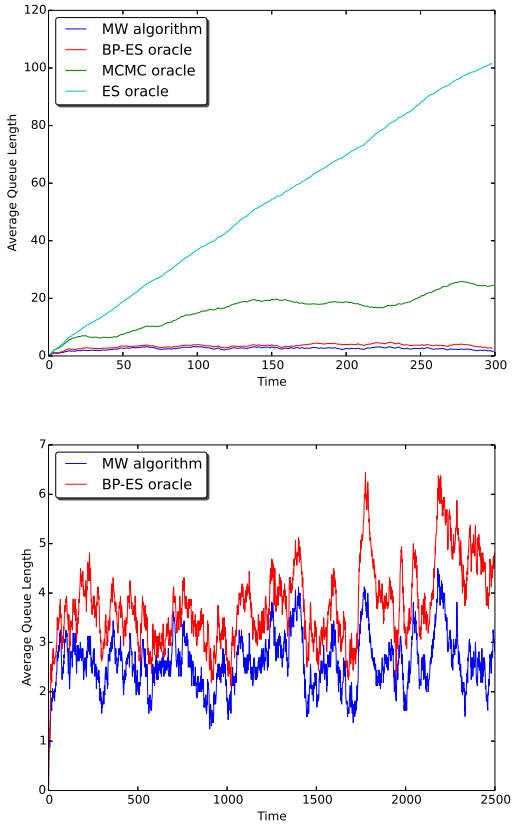

Figure 5: Performance of scheduling algorithms using MCMC, BP-ES and ES oracle systems in wireless network on 4×4 grid topology.

Wireless networks in grid topology. We also consider a wireless network with two dimensional grid interference topology. Due to the high complexity of the MW algorithm, we chose the 4×4 grid, i.e., 16 queues, where Figure 4.2 represents the 3×3 grid. We again compare the evolution of the average queue lengths of the MW algorithm and our scheduling algorithms using ES, MCMC and BP-ES oracle systems. We note that we consider a minor variant of BP-ES oracle so that it chooses a schedule greedily based on BP beliefs $[W_i - \sum_{j\in\mathcal{N}(i)} m_{j\to i}]$ similarly as the second BP

oracle in the input-queued switch network. We use the same weight functions as in the input-queued switch case and our experimental results are in Figure 5. We again observe that the performance of our scheduling algorithm using BP-ES oracle system is as good as that of the MW algorithm, although it is much simpler to code and faster to run.

7. CONCLUSION

The problem of dynamic resource allocation among network users contending resources has long been the subject of significant research for last four decades. In this paper, we develop a generic framework for designing resource allocation algorithms of low-complexity and high-performance via connecting iterative optimization and low-complexity scheduling. Our work shows sufficient conditions on queue-length functions so that a queue-based scheduling algorithms is throughput-optimal. To our best knowledge, our result is the first that establishes a rigorous connection between iterative optimization methods and low-complexity scheduling algorithms, which we believe provides various future directions and new insights in both areas.

8. REFERENCES

[1] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37:1936-1948, 1992.

[2] L. Georgiadis, M. J. Neely, and L. Tassiulas. Resource allocation and cross-layer control in wireless networks. *Foundations and Trends in Networking*, 1:1-144, 2006.

[3] T. E. Anderson, S. S. Owicki, J. B. Saxe, and C. P. Thacker. High-speed switch scheduling for local-area networks. *ACM Transactions on Computer Systems*, 11(4):319-352, 1993.

[4] N. McKeown. iSLIP: a scheduling algorithm for input-queued switches. *IEEE Transaction on Networking*, 7(2):188-201, 1999.

[5] C. Joo, X. Lin, and N. B. Shroff. Understanding the capacity region of the greedy maximal scheduling algorithm in multi-hop wireless networks. *IEEE INFOCOM*, 2008.

[6] M. Leconte, J. Ni, and R. Srikant. Improved bounds on the throughput efficiency of greedy maximal scheduling in wireless networks. *ACM MOBIHOC*, 2009.

[7] L. A. Goldberg and P. D. MacKenzie. Analysis of practical backoff protocols for contention resolution with multiple servers. *ACM SODA*, 1996.

[8] P. Gupta and A. L. Stolyar. Optimal throughput allocation in general randomaccess networks. ALLERTON Conference, 2006.

[9] P. Marbach, A. Eryilmaz, and A. Ozdaglar. Achievable rate region of csma schedulers in wireless networks with primary interference constraints. *IEEE Conference on Decision and Control*, 2007.

[10] S. Rajagopalan, D. Shah, and J. Shin. Network adiabatic theorem: An efficient randomized protocol for contention resolution. *ACM SIGMETRICS*, 2009.

[11] L. Jiang, D. Shah, J. Shin, and J. Walrand. Distributed random access algorithm: Scheduling and congestion control. *IEEE Transactions on Information Theory*, 56(12):6182-6207, 2010.

[12] D. Shah and J. Shin. Randomized scheduling algorithm for queueing networks. *Annals of Applied Probability*, 22:128-171, 2012.

[13] D. Shah, D. N. C. Tse, and J. N. Tsitsiklis. Hardness of low delay network scheduling. *IEEE Transactions on Information Theory*, 57(12):7810-7817, 2011.

[14] N. Bouman, S. Borst, and J. Leeuwaarden. Delays and mixing times in random-access networks. P *ACM SIGMETRICS*, 2013.

[15] D. Shah and J. Shin, Delay Optimal Queue-based CSMA. *ACM SIGMETRICS Performance Evaluation Review*, 38(1):373-374, 2010.

[16] M. Lotfinezhad and P. Marbach. Throughput-optimal Random Access with Order-Optimal Delay. *IEEE INFOCOM*, 2011.

[17] L. Jiang, M. Leconte, J. Ni, R. Srikant, and J. Walrand. Fast Mixing of Parallel Glauber Dynamics and Low-Delay CSMA Scheduling. *IEEE INFOCOM*, 2011.

[18] P. K. Huang and X. Lin, Improving the Delay Performance of CSMA Algorithms: A Virtual Multi-Channel Approach. *IEEE INFOCOM*, 2013.

[19] N. Abramson and F. Kuo (Editors). The aloha system. *Computer- Communication Networks*, 1973.

[20] J. G. Dai and B. Prabhakar. The throughput of data switches with and without speedup. *IEEE INFOCOM*, 2000.

[21] S. Kumar, P. Giaccone, and E. Leonardi. Rate stability of stable-marriage scheduling algorithms in input-queued switches. ALLERTON Conference, 2002.

[22] L. Tassiulas. Linear complexity algorithms for maximum throughput in radio networks and input queued switches. *INFOCOM*, 1998.

[23] P. Giaccone, B. Prabhakar, and D. Shah. Randomized scheduling algorithms for high-aggregate bandwidth switches. *IEEE Journal on Selected Areas in Communications High-performance electronic switches/routers for high-speed internet*, 21(4):546-559, 2003.

[24] E. Modiano, D. Shah, and G. Zussman. Maximizing throughput in wireless network via gossiping. *ACM SIGMETRICS*, 2006.

[25] A. Dimakis and J. Walrand. Sufficient conditions for stability of longest-queue- first scheduling: second-order properties using fluid limits. *Advances in Applied Probability*, 38(2):505-521, 2006.

[26] D. A. Levin, Y. Peres, and E. L. Wilmer. Markov Chains and Mixing Times. *American Mathematical Society*, 2008.

[27] J. Yedidia, W. Freeman, and Y. Weiss. Constructing free energy approximations and generalized belief propagation algorithms. *IEEE Transactions on Information Theory*, 51:2282-2312, 2004.

[28] G. Cornuejols. Valid Inequalities for Mixed Integer Linear Programs. Mathematical Programming Series B, 112:3-44, 2008.

[29] S. Foss and T. Konstantopoulos. An overview of some stochastic stability methods. *Journal of Operations Research, Society of Japan*, 47(4), 2004.

[30] M. I. Jordan. Graphical Models. *Statistical Science*, 19:140-155, 2004.

[31] M. Bayati, D. Shah, and M. Sharma. Max-product for maximum weight matching: Convergence, correctness, and lp duality. *IEEE Transactions on Information Theory*, 54(3): 1241-1251, 2008.

[32] S. Sanghavi, D. Malioutov, and A. S. Willsky. Linear programming analysis of loopy belief propagation for weighted matching. *NIPS*, 2007.

[33] S. Atalla, D. Cuda, P. Giaccone, M. Pretti. Belief-propagation assisted scheduling in input-queued switches. *IEEE 18th Annual Symposium on High Performance Interconnects (HOTI)*, 2010

[34] D. Shah. Gossip algorithms. *Now Publishers Inc*, 2009.

APPENDIX

A. PROOFS OF AUXILIARY INEQUALITIES

Here, we provide the proofs of (14)-(17) for large enough x. Due to conditions **C2** and **C5**,

$$\lim_{x\to\infty} \frac{2n}{f\left((1-c)x\right)} = 0,$$
$$\lim_{x\to\infty} \frac{h(f((1+c)x+2))}{x} = 0,$$

which proves (14) and (16).

Now, (16) follows from the condition **C1** and **C2**:

$$\lim_{x\to\infty} \frac{g((1+c)x)}{f((1-c)x)} \le \lim_{x\to\infty} \frac{2}{1-c}\frac{g((1+c)x)}{f(1+c)x} = 0,$$

where we use the following property for concave function f:

$$\frac{f((1+c)x)}{f((1-c)x)} \le \frac{1}{1-c}\frac{f((1+c)x)}{f(2x)} \le \frac{2}{1-c}.$$

Finally, for (17), let $x' = x + 2/(1+c) > x$ and we observe that

$$\begin{aligned}
& f'(f^{-1}(g((1-c)x)))\, h(f((1+c)x+2)) \\
&= f'(f^{-1}(g((1-c)x)))\, h(f((1+c)x')) \\
&= \frac{f'(f^{-1}(g((1-c)x)))}{f'(f^{-1}(g((1-c)x')))} \\
&\quad \times f'(f^{-1}(g((1-c)x')))\, h(f((1+c)x')) \\
&\le f'(f^{-1}(g((1-c)x')))\, h(f((1+c)x')) \to 0,
\end{aligned}$$

as $x \to \infty$ due to the condition **C6**. This completes the proof of (17).